Fine Tuning

THE ELECTRO-FLUIDIC AUTOPILOT

By Doug Garner (EAA 15611)
MS494 NASA Langley
Hampton, VA 23665
(804) 827-3321

History, Philosophy and Stuff

It has now been nearly seven years since we flew the first electro-fluidic autopilot here at NASA Langley. My first article on this subject (Ref. 1) simply described the project without giving any construction details; I had not yet come to think of this as a homebuilder's project. In a second article (Ref. 2), I tried to provide just enough technical information to allow the intrepid tinker to get started in this field, and to produce some working equipment. I knew that some sticky problems remained, and I began to get a little concerned to find people happily building wing levelers from the meager information that had been published, putting them in their airplanes, and asking for construction information on the magnetic heading reference add-on that makes the wing leveler into an honest autopilot.

Well, you can't make a good autopilot out of a bad wing leveler just by adding a heading reference, so I figured it was about time to go back and clean up our act and give you enough information to build a professional quality wing leveler before I got into a magnetometer construction article.

As with most homebuilt projects, it is fairly easy to make an electro-fluidic autopilot "that works", but it takes some extra thought and effort to include the design features and make the adjustments necessary to produce a really useful device that does not require continuous tinkering and fiddling to keep it in operation. If you enjoy this sort of diversion, read no further. If you want an instrument that will do its job with a minimum of fuss, and are willing to invest the time and effort required to understand the more subtle goings-on in this device, and make the appropriate adjustments, the following discussion should help.

One major problem encountered by the amateur autopilot builder has been that of zero shift of the output voltage from the rate sensor as a result of changes in air temperature, density and flow rate. A certain amount of drift can be overlooked in the basic wing leveler, since you need to adjust the trim knob occasionally to compensate for changes in lateral trim of the airplane anyhow, but if you are driving a turn coordinator indicator or adding a fully compensated magnetometer heading reference, it doesn't take much drift to spoil your fun.

Although the nature of the thermal pickoff used in the rate sensor and the lack of uniformity in the thermistors available for these pickoffs will limit what can be achieved in this area, we can improve the situation by an order of magnitude or so by paying attention to the details of design and adjustment of the rate sensor and its circuitry.

At least one major item remains to be dealt with, and that is the question of a satisfactory servo for this system. We are currently addressing this problem, and will try to give you all the answers as soon as we are really sure we know what we are talking about.

In this article, I shall try to cover the drift problem in some detail, bring you up to date on current technology, fill in some gaps, and tell you about some of the tricks and gimmicks we've learned over the last seven years. Another article being prepared concurrently with this will tell you how to build a simple magnetometer heading reference that will convert your electro-fluidic wing leveler into an honest, two-axis autopilot that will hold a set course all day long without attention.

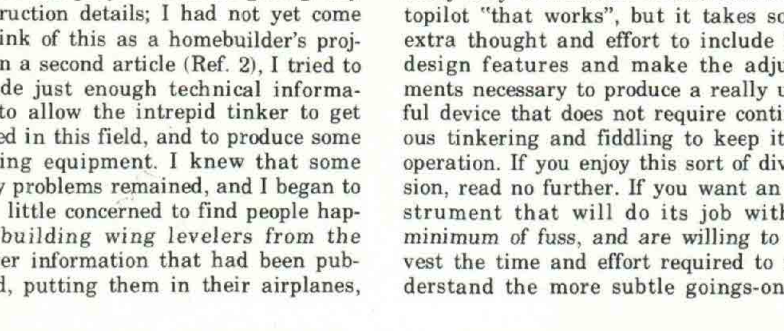

FIGURE 1
THE BASIC RATE SENSOR CIRCUIT

Rate Sensor Operation

In order to address the drift problem intelligently, we must get into some of the grimy details of operation of the rate sensor. The Laminar Jet Rate Sensor, which forms the heart of the simple wing-leveler autopilots described in Ref. 1 and 2, depends upon the inertial deflection of a jet of air to sense rotation of the aircraft about its roll and yaw axes. The particular version of this rate sensor which we use (Ref. 3), detects the deflection of the jet by its cooling effect on a pair of self-heated thermistors, positioned at either side of the jet's path. (A thermistor is a resistor compounded from an esoteric mixture of metallic oxides and floor sweepings from select hen houses. Its salient feature is its large variation of electrical resistance with temperature.)

Each thermistor forms one leg of a separate Wheatstone bridge circuit (Fig. 1) which is in balance only when the thermistor is at a specific temperature (about 150 degrees F., in this case). The output of each bridge is connected to an operational amplifier (a high-gain DC amplifier) whose function is to supply just enough current through the bridge to heat its thermistor to a temperature sufficient to balance the bridge.

As the jet is deflected to blow more on one thermistor and less on the other, it requires more electrical power to maintain the first thermistor at its required operating temperature and less for the second, so the voltage from the first bridge amplifier goes up and that from the second goes down. This **difference** in voltage between the two bridge amplifiers is then seen to be a measure of the jet deflection and, therefore, the rate of rotation of the aircraft in inertial space. With all the nonlinearities inherent in this sort of an operation, you would think this voltage difference would be anything but a linear function of turn rate, but, as we will see later, there are some nice, wide, linear regions in the input/output relation, which we can exploit by choosing the right thermistor spacing.

O.K., so this thing works by sensing the difference in cooling effect on the two self-heated thermistors, produced by the inertial deflection of the air jet, which causes it to blow more on one thermistor and less on the other. What we are really sensing is the difference in thermal dissipation of the two hot thermistors (the difference in the rate of heat flow from each thermistor to its environment), and, unfortunately, a number of other factors influence this heat flow much more than does the deflection of the air jet.

Two major contributors are the temperature and the density of the air in the sensor, and to give you an idea of the importance of this problem, the change in bridge voltage for a change in air temperature of one degree Fahrenheit will be more than twenty times as great as that for a one-degree-per-second rate of turn.

Common Mode Rejection

Fortunately, most of the things that effect the thermal dissipation of the thermistors tend to effect both thermistors in the same way, so that both bridge voltages will change by the same amount, whereas the jet-deflection effect, that we are really interested in, causes one bridge voltage to go up and the other to go down. This allows us to use an item from the instrument engineer's bag of tricks called "Common Mode Rejection",

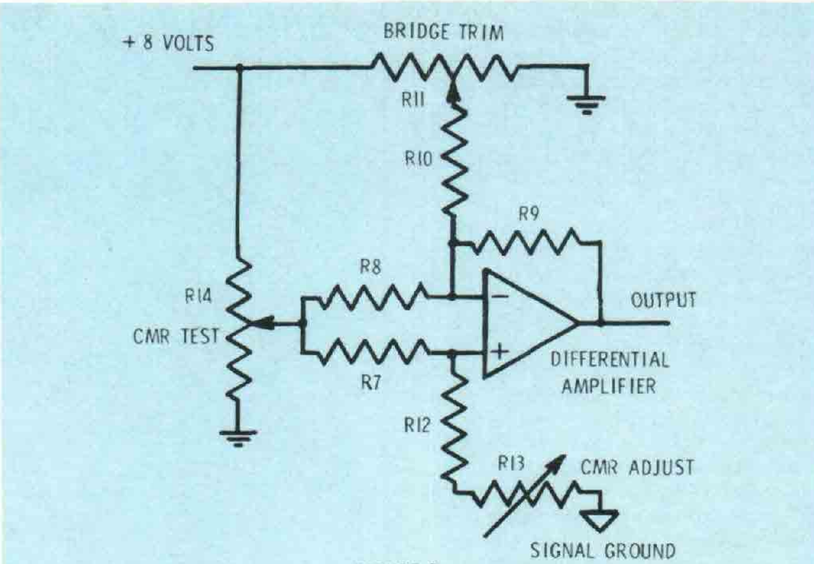

FIGUPE 2
THE COMMON MODE REJECTION ADJUSTMENT

that lets us ignore equal voltage changes and see only the **difference** between the two bridge voltages.

In Figure 1 we see that the outputs of the two bridge amplifiers are fed to the two inputs of a third operational amplifier called the "differential amplifier", whose function is to subtract one bridge amplifier voltage from the other and to amplify only the **difference** between these two voltages. In order to accomplish this in a satisfactory manner, the gains of the differential amplifier must be precisely equal for both inputs, but must be positive for one input and negative for the other. We are not too concerned about the exact values of these gains, but we are most interested in having the gain associated with the positive input precisely the same as that of the negative input; otherwise, when the two inputs increase or decrease by exactly the same amount, as they should in response to changes in air temperature, etc., we will get an unwanted change in output.

There are a number of different circuits for attaining this laudable goal, and these are called "instrument amplifiers". Any good book on op-amp technology, like Ref. 4, will supply more theory on this subject than most of us can comfortably assimilate. The circuit shown in Figure 1 is about as simple as these things get, but it takes just the right combination of input and feedback resistors to give us perfect common mode rejection. The voltage gain for the negative input will be equal to the value of R9 divided by that of R8. This gain is set to saturate the following limiter amplifiers at a signal voltage corresponding to a rate of about plus and minus five degrees per second at the rate sensor.

For reasons too gruesome to contemplate in this article, the gain of the positive input of an op-amp in this configuration will always want to be a bit higher than that of the negative input, so the trick is to add just enough shunt resistance at the positive input to make the two gains come out equal. The value of the shunt resistor can be calculated, but it is easier and more accurate to arrive at this value experimentally.

Before we get into this procedure, there are a couple of other small items to clear up. First, in spite of all our efforts to make the two thermistor bridges identical and to balance up the rate sensor mechanically, we will end up with some difference in the two bridge amplifier output voltages at zero rate input. The bridge trim potentiometer R11 and its summing resistor R10 will allow us to trim out this difference. This added shunt on the negative input of the amplifier will increase the gain of the positive input even more, so the value of the bridge trim potentiometer is made as small as practical with respect to R10, so that the bridge trim adjustment will have a minimum effect on the final common-mode-rejection adjustment.

The other matter concerns the "signal ground". We are using a single, regulated, 8 volt power supply for most of the signal processing amplifiers, and in order to handle both positive and negative signals (turns to the right and turns to the left) we have to establish our zero signal value at approximately 4 volts, so that positive signals can go above signal ground, and negative signals below by an equal amount. Actually, we set it at about 3.6 volts to keep it near the center of the output voltage swing of these particular opamps, operating on the eight volt power supply. The symbol for the signal ground is a triangle, while that for the power ground is the conventional ground symbol.

Getting back to that all-important common-mode-rejection adjustment, one simple way to make this is to disconnect the two differential amplifier inputs from the thermistor bridges and connect them both to a variable voltage input (R14) as in Figure 2. As this input voltage is varied up and down, R13 is adjusted to give zero **change** in output voltage from the differential amplifier. What we are doing is adjusting the shunt resistor (R12 plus R13), so that the differential amplifier has a gain of absolutely zero for equal changes in both input voltages, and we are making sure we have exactly equal changes in both input voltages by connecting both inputs together for the test.

The complete circuit in Figure 3 has provisions for this adjustment built in. The switch, SW1, makes the proper circuit changes for the test, and R14 provides the variable voltage input.

In making this adjustment, we have to look at the voltage output from the differential amplifier. For convenience, I use a zero-center voltmeter with a full-scale range of plus to minus four volts, connected between signal ground and the output of the amplifier. If you just have an ordinary, garden-variety voltmeter, set it for eight volts or more full scale, and hook the negative lead to the power supply ground. First connect the positive lead to signal ground, and mark this reading (about 3.6 volts) with a bit of tape on the meter glass to indicate signal

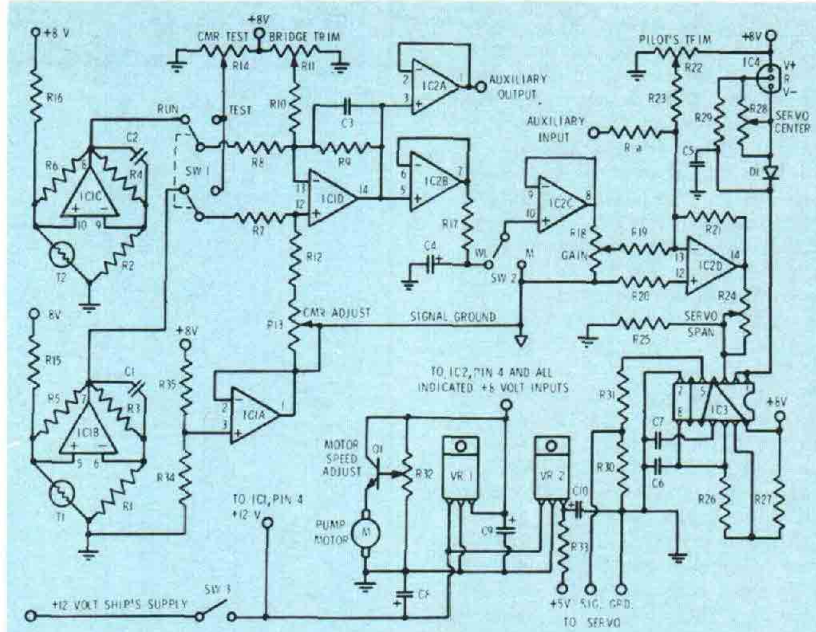

FIGURE 3
THE COMPLETE WING LEVELER CIRCUIT

PARTS LIST

ALL FIXED RESISTORS 1/4 WATT EXCEPT R33

R1, R2	10,000 ohms - Matched Pair
R3, R4	4,700 ohms - Matched Pair
R5, R6	220 ohms - Matched Pair
R7, R8, R30	1000 ohms
R9	27,000 ohms
R10, R19, R20	10,000 ohms
R11	1,000 ohm trim pot
R12	4,700 ohms
R13, R24, R28, R32	10,000 ohm trim pots
R14	5,000 ohm trim pot
R15, R16, R17	1.0 megohm
R18, R22	10,000 ohm volume control pots with knobs - linear taper
R21, R23	100,000 ohms
R25	5,600 ohms
R26	3,300 ohms
R27	270,000 ohms
R29	47,000 ohms
R31	270 ohms
R33	1 to 3 ohms, 1 watt (See Text)
R34	1,200 ohms
R35	1,500 ohms
R a	Select for input
C1, C2, C3, C6	0.1 mfd Mylar
C4	1.0 mfd, 35 WV Tantalum
C5, C7	0.01 mfd Mylar
C8, C9, C10	10 mfd, 25 WV Electrolytic
IC1, IC2	LM324 (RS 276-1711) Quad Op Amp (Pin 11 = Grd)
IC3	LM556 (RS 276-1728) Dual Timer
IC4	LM 334 (RS 276-1734) Current Source
VR1	LM340-8 Voltage Regulator 8 Volts
VR2	LM340-5 Voltage Regulator 5 Volts
Q 1	2N4401 (See Text)
D 1	IN914 (RS 276-1122) EPOXY TO IC4
M	Motor - Micro-Mo 050/B04 (See Text)
T1, T2	Thermistors - Fenwal GB32L1

zero; then connect the positive lead to the amplifier output to make your readings.

This adjustment can be confusing, but I have found the following routine helps: with SW1 in the test position, and R13 at one end of its travel, the voltmeter connected to the output of the differential amplifier (pin 14, IC1D) will move in the same direction as R14. With R13 at the other extreme, the voltmeter will move in the opposite direction from R14. By remembering which way is which, you can keep resetting R13 until you get no motion at all from the voltmeter needle as R14 is rotated from one end to the other. Reset the output of the amplifier to near signal zero with R11 each time you make a change to R13. The "Auxiliary Output" terminal of Figure 3 is handier to get at than pin 14 of IC1D, and will give the same results. You must, of course, switch back to the "run" position to get back in business. This adjustment should only have to be made once.

I have belabored this common-mode-rejection adjustment rather heavily because it is important, and because it seems to be rather difficult to get across to many people. I hope I have not bored the rest of you too deeply, and that you will take the time to do this right. It will do wonders for the zero stability of your wing leveler.

A More Sophisticated Circuit

Having introduced Figure 3 in the preceding explanation, this is probably a good time to examine it in more detail. It is basically the circuit shown in Reference 2, but contains some minor modifications to allow for the common-mode-rejection adjustment already discussed, and to avoid the interaction between gain and trim adjustments encountered

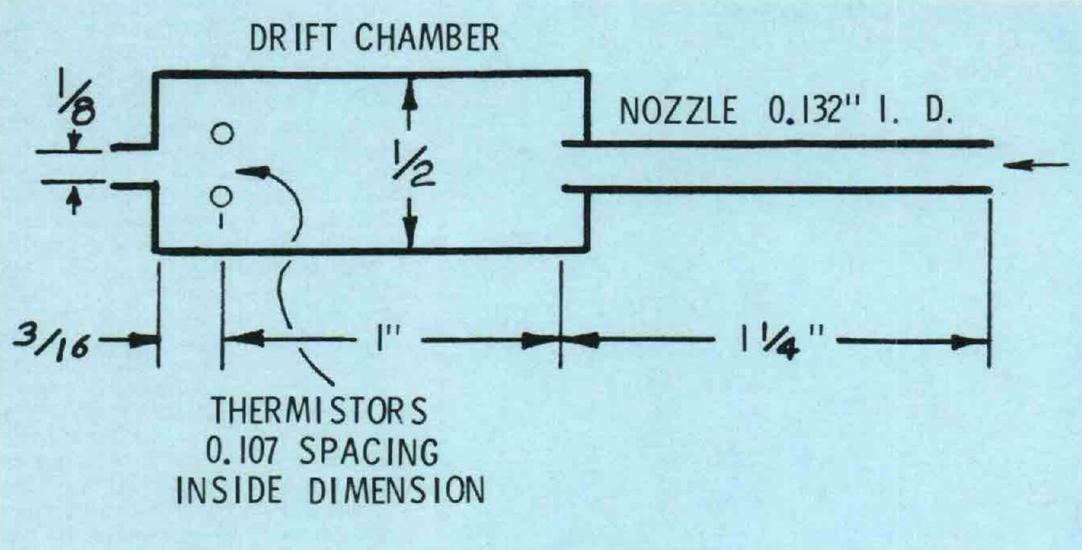

BASIC DIMENSIONS OF THE RATE SENSOR
FIGURE 4

in the previous circuit. It also allows the rate sensor to operate at considerably lower temperatures without saturating the bridge and differential amplifiers and provides a bit more linear operation of the servo.

The most obvious change is the addition of another op-amp integrated circuit chip. IC1 and IC2 are both "quad op-amps", which means that each unit, or "chip", contains four operational amplifiers, which are completely independent of one another except for their power supply connections.

In order to allow a greater range of output voltages from the bridge amplifiers, so that they can cope with lower environmental temperatures, we now power these amplifiers from the unregulated, 12 volt power supply. This may sound like a bad move, but these great little amplifiers are so designed that reasonable variations in the power supply voltage have virtually no effect on the signal level. The differential amplifier must also be operated on the 12 volt supply, so that its inputs will not be saturated by the higher output voltages from the bridge amplifiers at low temperatures, so it goes on the same chip as the bridge amplifiers.

You will notice a few components in the bridge circuits of Figure 3 that did not appear in Figure 1. These op-amps have very high gains, and, as used in these bridge circuits, have a lot of positive feedback at the higher frequencies, so they are prone to burst into wild oscillations. C1 and C2 add extra negative feedback at the higher frequencies, and prevent oscillation.

We have a chicken-and-egg situation with these bridge circuits. There must be some voltage on the bridge before the amplifier can sense the fact that it needs to apply voltage to the bridge, so it may never get off ground zero when the power is first applied. R15 and R16 give the amplifier a nudge in the right direction to get things started off reliably.

Once the two bridge amplifier outputs have been combined into one signal voltage by the differential amplifier, we would like this signal voltage to be symmetrically limited on either side of the "signal ground" level. The easiest way to accomplish this is to arrange for the following amplifiers to saturate at equal voltages above and below signal ground, so we operate the remaining amplifiers in the string (all on IC2) on regulated 8 volts, and set the signal-ground voltage level halfway between the upper and lower saturation voltage levels of these amplifiers as noted previously.

IC1A supplies a low impedance source for the "signal ground" voltage, determined by the voltage divider made up of R34 and R35.

IC2A, IC2B and IC2C are all connected as voltage followers, meaning that they have a gain of exactly one, and that they do not invert the polarity of the signal. IC2B simply acts as a limiter on the signal voltage before it is applied to the lag circuit, consisting of R17 and C4.

As explained in Reference 1, this lag circuit contributes somewhat to the stability of the aircraft-wing-leveler combination, and, more importantly, it keeps the servo from responding to all the little high frequency gust disturbances, and growing old before its time. This lag circuit can be used only with aileron control. It will have a destabilizing effect if a rudder servo is to be used. It is a good idea to leave C4 out of the circuit until you get the rest of the system operating properly, as the slow response of the system with the lag circuit in operation can be confusing when trying to make other adjustments.

SW2 selects either the wing-leveler or the manual-trim mode of operation. In the manual-trim mode (M), the servo can be positioned by the pilot's trim control (R22), but the wing-leveler function is inoperative. Since the warm-up cycle of the rate sensor, lasting about six seconds, features hard-over signals first one way and then the other, it is considerably less exciting to put SW2 in the manual-trim position beftre switching on the power (SW3) in flight.

IC2C acts as a high-input-impedance buffer to avoid loading the lag circuit, and drives the wing-leveler-gain potentiometer, R18. It is a matter of choice whether this is an internal trim pot or is brought out where the pilot can adjust it in flight. I generally work hard to reduce the number of knobs and meters and doodads on the control panel to a minimum, but, since the sensitivity of the rate sensor in this set-up changes with air density (it goes down by about one-half at ten thousand feet), and the response of the airplane also changes with flight conditions, it is probably good to be able to get hold of this one, at least until you get used to the way the system performs.

The gain control, as shown, covers a rather wide range, so, if you opt to leave it on the panel, you may want to add some end resistors to limit its range to the range of gains you find useful in flight. There's nothing much worse than a "touchy" adjustment on a flight panel.

IC2D acts as a summing amplifier for the rate signal, the pilot's trim control and whatever you choose to put into the auxiliary input (a magnetic heading reference, radio navigation aid, turn command switch, etc.). It also adds a gain of ten, to get the overall wing-leveler gain up to a maximum of full servo travel for 0.5 degrees per second rate of turn.

IC2A is used as a buffer to supply a rate signal voltage to a turn coordinator indicator, the compensation coil of a magnetic heading reference or whatever.

We could have gotten by nicely with a few less op-amps, but our requirement for operation on two different supply voltages bought us a minimum of two op-

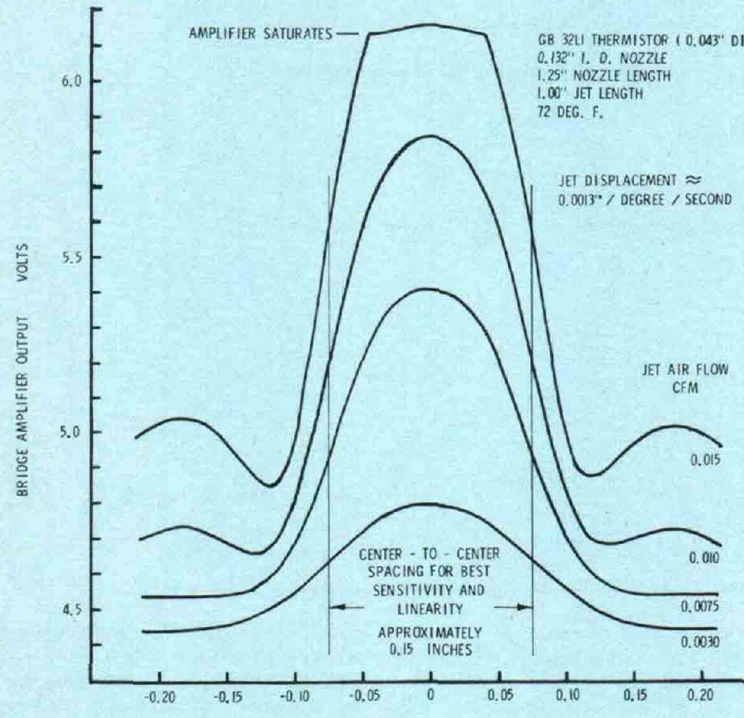

FIGURE 5
RATE SENSOR PERFORMANCE CHARACTERISTICS

amp chips with four amplifiers on each, so we might as well live it up a little.

About Toy Servos

IC3 converts the D C signal voltage from the summing amplifier into the weird pulse-width-modulated signal that Radio-Control model servos understand.

These are evil little mechanisms, and ill suited for this job, but they are cheap and readily available, and they can be made to work, so we shall confine this discussion to a few minor circuit features which may help to make the best of the situation. Work is currently in progress on more practical servo designs. More about this latter, perhaps.

The pulse coding circuit shown in Figure 3 is virtually the same as that in Reference 2, which was cobbled up in haste from a couple of cook book circuits to get a demonstration model going for the '77 Oshkosh forum. A few minor changes have been made to improve linearity and to make initial adjustment less tedious.

The signal required for the typical R C servo (I used a Heathkit Model GDA-1205-8) is a positive pulse of about five volts amplitude, with a repetition rate of about 60 per second (I have found that I could vary the pulse rate from 40 to 250 per second without noticeable effect on servo operation). The position of the servo shaft is directly proportional to the width, or duration, of this pulse, from one to two milliseconds, with the center position corresponding to 1.5 milliseconds.

The secret decoder circuit that changes pulse width into servo position lies buried in a proprietary integrated circuit chip in the servo, and is more trouble to modify than the project is worth, so, if you must use R C servos, you are stuck with generating this dumb pulse-width signal.

Experience has shown that a detailed explanation of the operation of the pulse coding circuit implemented by IC3 takes several closely typed pages and is excruciatingly boring. Suffice it to say that the pulse repetition rate is determined by R26, R27 and C6, and the basic pulse width by C5 and its charging circuitry. Upon sober reflection, R26, R27 and C6 have been changed to more suitable values than those shown in Reference 2. The basic pulse width is modified by the signal voltage on pin 3, to give us the signal-voltage-to-pulse-width conversion we are after. The resultant pulse train comes out of pin 5, and is reduced to a 5 volt level by the voltage divider, R30, R31, so as not to overdrive the decoder in the servo.

In the original circuit of Reference 2, C5 was chareed through a resistor, resulting in a somewhat non-linear charging rate which was reflected in a similar nonlinear relation between signal voltage and pulse width. This was no big deal, amounting to a difference of about ten percent between the travel of the servo in one direction and that in the other, and should not effect the overall performance of the system, but this kind of thing upsets the decorum of the purist.

Figure 3 shows you how to do it right, by replacing the charging resistor with a constant-current source. IC4, which will charge C5 at a nice, linear rate. Unfortunately, this thing is temperature sensitive, and must be compensated by the addition of D1 and R29. D1 is stuck onto the flat side of IC4 with a dab of epoxy to help keep them at the same temperature. The amount of compensation is determined by R29, and if you really expect a hot summer, you might check this out by warming up IC4 with a hair dryer (don't melt it). If the servo position shifts with temperature, adjust the value of R29 until it doesn't.

If all this seems like more trouble than it's worth, just replace IC4, D1, R28 and R29 with a 500,000 ohm trim pot and take your licks. This, at least, is not temperature sensitive.

Disk type ceramic capacitors are quite temperature sensitive, and should not be used for C5. Invest in a good mylar capacitor here. It only costs a few cents extra to go first class with the other small capacitors, as well.

In the original circuit, the adjustments for servo centering and maximum displacement were very interactive, and getting these both set up correctly at the same time was pretty tedious. The normal voltage level at pin 3 of IC3 is set by an internal voltage divider at two-thirds of supply voltage, while our "signal zero" voltage is a bit less than one-half supply voltage. The addition of R25 loads the voltage divider in IC3 down to about "signal zero" level, and gets both circuits started at the same level.

To make the servo zero and span adjustments, set SW2 to the "M" position, and adjust R22 for zero voltage between pin 14, IC2D and signal ground. Then center the servo with R28. Set R22 to one end of its travel, and set the maximum desired servo travel with R24. If the constant current version (IC4) is used, the servo should then have equal travel in each direction. It is very important to set the maximum servo travel low enough that it never strikes its mechanical stops, if you want the thing to survive for any length of time.

VR1 is an eight-volt voltage regulator, which supplies power to most of the circuits, with the exception of IC1 and the servo. I chose this voltage because it is about as high as you can go and still maintain proper operation of the regulator over the full range of voltages normally encountered in an aircraft's power system. VR2 supplies five volts of power to the R C servo because that is the voltage at which it is rated. These servos draw peak currents around one amp, which is the maximum current rating of the usual voltage regulator chip.

These peaks can trip the current limiting circuitry built into the regulator chip and lead to a constant "jitter" of the servo, with a resulting high standby current. There are VR chips that go a bit higher in current rating, but they are bulky, expensive and hard to find. My solution to all this was to add a current-limiting resistor of one to three ohms and one-watt capacity (R33). If at all practical, this resistor and VR2 should be mounted on a separate heat sink (a piece of one-eighth aluminum about two inches square) outside the wing-leveler case, as these components produce considerable heat and add to your temperature drift problems. The rest of the circuitry draws less than one-tenth of an amp, so you can mount V 1 right on the circuit board. C8, C9 and C10 are needed for stable operation of the regulators.

R32 and Q1 form a low-impedance voltage source to drive the air pump motor, if one is used, at constant speed. Almost any NPN silicon transistor should do for Q1, if it has a current rating adequate for the pump motor you choose.

Bridge Components

Making the common mode rejection adjustment to the differential amplifier is the easy half of the solution to the drift problem. The tough part is in the selection of components for the two bridges so that both bridges will, indeed, track perfectly over a wide range of ambient temperatures and densities. You can use precision resistors for the fixed resistors in the bridges, or you can go through a small batch of ordinary one-quarter watt carbon or film resistors of each value with a good ohmmeter and pick out the best matched pairs (R1-R2, R3-R4, R5-R6).

But the thermistors are something else again. Even if the fixed resistors are perfectly matched, the two thermistors will not operate at the same temperature unless they happen to have the same temperature/resistance characteristics at the desired operating temperature; and if they operate at different temperatures, their coefficients of thermal dissipation will not change by equal amounts with changes in ambient temperature and pressure. Differences in size and surface area and the proximity of the various parts of the mounting structure also have their effects.

All this means that we can't really hope to select a perfectly matched pair of thermistors, but there are some things we can do to improve the situation considerably.

The thermistors we used in the original research autopilots were simply matched as closely as possible for resistance at room temperature. This room-temperature resistance varies rather widely for these small, bead thermistors, but the percentage change in resistance at other temperatures seems quite consistent, so if you match them at room temperature, they will probably be about right at operating temperature. We later tried some pretty complicated ways of matching thermal dissipation constants at operating temperatures, but without much improvement.

Of course, if you just want to buy one pair of thermistors to build one autopilot, you can't do much about the matching process, but you can purchase "matched pairs" from at least one source (Ref. 5). Although there is just no way to match up all the parameters, these will be a great improvement over random selection.

Both Humphrey, Inc. and Hamilton Standard now market sophisticated (and expensive) versions of the laminar jet rate sensor, and both use hot-wire pickoffs rather than hot thermistors. Two pieces of wire from the same spool should match up better than two thermistors, and I expect the wire pickoffs are much more consistent and easier to get along with than those made from thermistors. Anyone with the requisite skills and equipment to fabricate a hot-wire pickoff would not be wasting his time giving it a try.

Rate Sensor Design

Having done the best you can on thermistor and bridge resistor selection, it will be time to worry about the mechanical arrangement of the various parts of the rate sensor. The basic dimensions shown in Figure 4 have worked out very well and should not be changed, unless you have a pretty good reason (you will note some minor variations from the dimensions given in Ref. 2). The nozzle must be about this long to establish a good, stable laminar-flow jet in the noise and vibration environment of the aircraft. The distance from the nozzle exit to the thermistors is a compromise between sensitivity and jet stability — the longer the distance, the more sensitive the rate sensor; but the jet will only stay laminar so far out.

Nozzle diameter is not very critical, but it is closely related to thermistor spacing. Figure 5 shows some data we took to determine the optimum spacing for the thermistors in a rate sensor with the dimensions of Figure 4. Here we traversed a single thermistor across the air jet in the plane of the thermistor mount and plotted its displacement against the bridge voltage it produced for various jet flow rates. You can see that the curves have rather long, straight sections at either side, and that these straight sections have the greatest slope of any part of the curve. Now if we space our thermistors so that they fall at the centers of these straight sections, we are assured of maximum sensitivity and a linear relation of output voltage to turn rate. The ideal center-to-center thermistor spacing for this nozzle diameter is about 0.150 inches, or an inside spacing of 0.107 inches. The shank of a number 37 drill is about right as a gage.

A few thousandths one way or the other on either the thermistor spacing or the nozzle diameter will not make much difference, as the jet deflection is only a bit over one thousandth of an inch per degree per second.

We used this particular nozzle diameter, as it is the inside diameter of stock 3/16 O.D. brass, aluminum and copper tubing. If you insist on using another tube diameter, you will have to run your own data or make your own guesses. The big danger in guessing at the thermistor spacing is that you could easily end up near one end of the linear region and have something that works great on the bench, but will shift into the low gain, nonlinear region if operating conditions change slightly later on.

You will notice that the upper curve is chopped off rather abruptly. This is due to saturation of the bridge amplifier, which was being operated from an eight-volt power supply for this test. As ambient temperature is reduced, all the curves will move upward on the bridge voltage scale, until saturation occurs at normal flow rates and turn rates at about 30 degrees F. The strategy of operating the bridge and differential amplifiers on twelve volts, as already noted, will bring the lower temperature limit down into brass monkey territory.

Precise centering of the air jet with respect to the thermistors is also important if you want your rate sensor to remain stable under flight conditions. You can center things up to eyeball accuracy by peeking down the nozzle and centering up the thermistors so that you can see equal portions of each one. This is the place to start and to get things into operation, but don't stop there. I gave some detailed instructions on this centering business in Reference 2, but since then I have seen some rate sensors built in such a way that no adjustment was possible, so we will go through it again here.

With the jet off to one side a bit, we can still bring the differential amplifier back to zero voltage output for zero rate input with the electrical bridge trim adjustment (R11). It should be obvious, however, that under these conditions, changes in flow rate, density and temperature of the air will not affect both thermistors equally, since one is in a higher flow region than the other at zero rate input.

Ideally, we should like to center the jet to the extent that varying the jet velocity over the full, usable range will not give us any zero shift (though it will, of course, give some change in sensitivity). Unfortunately, a number of other factors get into the act, and this is not really possible in practice, but it is well worth the effort to get as close as you can.

At one time, I showed a circuit in my forum notes that controlled the speed of the motor driving the jet air supply pump as a function of the sum of the two bridge output voltages. This worked to keep the total cooling effect of the jet on the two thermistors constant, and tended to avoid any zero shift or sensitivity changes due to changes in air velocity or density, even with a misaligned jet.

Using this pump-motor control circuit, I was able to take the rate sensor up to a pressure altitude of fifteen thousand feet without any zero shift or change in sensitivity, at which point the air pump motor had reached its maximum speed.

Unfortunately, air temperature also has a great effect on bridge voltage, with the result that the pump motor would speed up as air temperature increased, soon driving the jet into the turbulent-flow region on a hot day, or would stop completely, as the temperature was reduced. If this circuit were adequately compensated for ambient temperature, it might prove to be a useful asset, but in its present form, it's a potential disaster. If you have a copy of this circuit, please burn it!

Center That Jet

Getting back to the jet-centering procedure, you can use any sort of an arrangement that will allow a fine adjustment of the nozzle centerline with respect to the thermistors. We mount our thermistors on a flat framework made from printed circuit board material and move this thermistor mount with respect to the jet, using a special adjusting tool to be shown later.

The jet-centering procedure consists of adjusting the electrical bridge trim (R11, Figure 3) for zero output of the differential amplifier with no airflow from the jet, and then, with the jet in operation, adjusting the mechanical alignment of the jet to return the amplifier output to zero. As these two adjustments are somewhat interactive, it pays to go through the whole procedure several times, until things settle down.

With no air flowing through the nozzle, the thermistor assembly is very sensitive to acceleration. If it is positioned so that one thermistor is higher than the other, convection currents from the lower thermistor will rise and add heat to the upper one, producing an output signal (an inclinometer called a "Convectron", that worked like this was marketed shortly after World War II by Bendix, I believe).

This means that you have to position the rate sensor rather carefully during the above jet-centering procedure to minimize the effect of gravity on the no-flow adjustment. We mount the rate sensor on a small piece of aluminum angle which is clamped to the bench, so that both the center line of the nozzle and a line through the centers of both thermistors lie in a horizontal plane.

The Bringing-Up Ceremony

When you first apply power to your wing leveler (a traumatic event, at best), you should be able to make the no-flow zeroing adjustment, if everything is hooked up correctly. The rate sensor assembly is then unclamped from the bench, so that you can rotate it back and forth about a vertical axis while you slowly increase the air flow to the jet nozzle. At some point, the output voltage of the differential amplifier will start to respond to the rate input. This represents the minimum usable air flow. As the air flow is further increased (with the rate sensor stationary), you will reach a point where the output voltage begins to get very nervous, indicating the onset of turbulent flow. This is the maximum permissible flow rate. To play it safe, you should operate about halfway between these two flow rates. You will probably have to run the jet-centering procedure concurrently with this "bringing-up" routine, and you should, certainly, give it a final check after you determine the best flow rate for your sensor.

Air Supply

If you use the electrically driven air pump for your jet supply, you can adjust jet flow with R32. If you power the jet from the airplane's vacuum system, one of those little needle valves used in model airplane motors works fine as an adjustment, but you will need a vacuum regulator ahead of the needle valve to hold the vacuum constant with changes in engine speed. We cobbled up a simple vacuum regulator from a couple of pieces of plastic, a bit of rubberized fabric and a tire valve, that seems to do a good job. This should work, even with a venturi vacuum supply. The neat, built-in venturi shown in the KR Newsletter, Number 48, should take care of the rate sensor as well as a pair of vacuum aileron servos for KR and other foam-fiber glass builders.

Rate Sensor Construction

Figure 6 shows the mechanical details of the rate sensor we have developed here at the lab. This design uses the basic dimensions of Figure 4. We have tried to keep it as simple and easy to build and adjust as possible, since we still have to make our own. We are blessed with a lathe and a small vertical mill in our laboratory shop that make life a bit easier, but there is nothing in this design that cannot be handled with a drill press and some hand tools.

The large holes in the Plexiglass parts may be a problem to some. On the prototype, I drilled these with a one-half inch bit and then brought them up to size with one of those great little boring tools you see demonstrated at the EAA fly markets (Ref. 6). When drilling Plexiglass, use a vise or vise grips to hold small parts like these, and clear your drill of chips often. Large drill bits have a great tendency to jam in Plexiglass, and the spinning part makes an effective meat chopper. You can, of course, use other materials than those shown. Hardwood is not bad for the various rectangular blocks and the plug that holds the nozzle. This last item can be roughed out with a file, drilled and cemented to the nozzle tube, and then brought to size with a file, with the tube chucked in a drill press. Get these parts lined up as accurately as possible when you cement them together, or you will have trouble getting the thermistor assembly lined up with the jet later.

The nozzle must be made from a straight piece of tubing, and all internal burrs must be removed from each end to avoid disturbing the laminar flow of the jet. We use five-minute epoxy to glue the parts together. The end of the five-eighths tubing is set back one-sixteenth of an inch from the face of the block it is cemented into, to give added clearance for the globs of solder holding the thermistor leads in place. The inner end of the jet nozzle tube extends one-sixteenth of an inch beyond the face of the plug into which it is cemented.

The thermistor mount is made from glass-epoxy printed-circuit board with the copper cladding (one side only) etched as shown, to separate it into three segments. If you don't have the facilities to etch it, you can work the slots out with a hacksaw blade, making them all straight. For convenience, we shape the end to which the electrical connections are made to fit standard 0.156 inch P C board socket spacing, and make three-contact plugs for it by sawing up standard P C board sockets, like Radio Shack 276-1551.

Before soldering the thermistors in place, fill the etched or sawed slots with epoxy, and then sand this flat with fine sandpaper on a flat surface to avoid air leaks. Sand the two mating surfaces flat too, for the same reason.

The thermistor leads are passed through the #60 holes from the copper clad side, gently pulled tight from the other side and bent over to hold the thermistors in place until the leads can be soldered. Make sure the thermistors are centered in the three-eights inch diameter clearance hole. The leads are then soldered to the copper cladding, using as little solder as possible. Resin flux will not work with the platinumiridium thermistor leads. You will need to use one of the zinc chloride based liquid fluxes, such as Green Streak, to get the solder to tin these properly. Use a magnifying glass, and be sure the solder actually tins the thermistor leads. If you just glue them down with a glob of solder, without getting the solder to flow onto the leads, you will be sorry later.

These zinc chloride fluxes are the so-called "acid fluxes", and are generally deadly to electronic circuitry, so take extra care with cleanup. Get all the residual flux off by a gentle rinse in hot water, and a bit of scrubbing with the tip of a toothpick. While things are still hot, lay the assembly gently on a folded Kleenex, thermistor side down, to dry. Corrosion due to any remaining flux can be disastrous. Check the spacing with the shank end of a number 37 drill, as noted before. Don't touch the thermistors with your fingers — the grease you deposit on them will foul things up properly.

The thermistor mount pivots about the upper mounting screw for the centering adjustment. The adjusting tool shown in Figure 6 is soldered up from bits of drill rod, old nails or whatever. When making the centering adjustment, loosen the bottom screw completely and leave the top screw just a bit snug until the amount is properly adjusted; then tighten both screws, and check the adjustment once more.

Pump Motors

We use Micro-Mo 050/B04 motors (Ref. 7) to drive the little centrifugal air pump that draws air through the nozzle. These are rated at 12 volts, but, with the pump geometry shown, they supply the correct air flow at about three volts and draw about 2.5 milliamps. We have never had one fail. These are rather expensive, about $15.00, and a bit hard to get in small quantities. The motors supplied as replacements for R C servos are a fair second choice, but will draw a lot more current.

As a last resort, some of the small motors available from Radio Shack can

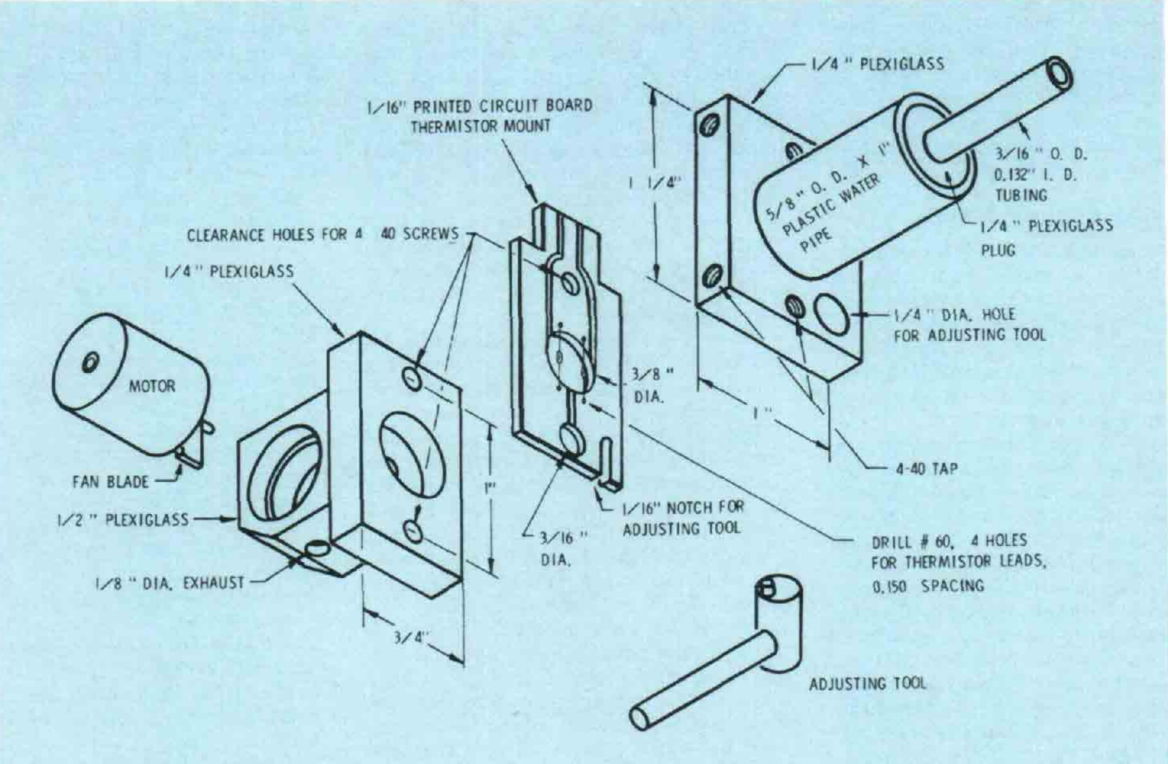

FIGURE 6
RATE SENSOR ASSEMBLY

be used, though they may require a huskier drive transistor (Q1), and won't have much of a service life. These vary greatly in quality and current requirements. Pick one with carbon brushes and the lowest current rating you can find.

Air Pump Design

Air pump construction is illustrated in Reference 2, and not so clearly in Figure 6 of this article. The outer casing is a one-half-inch hole in a block cemented to the rate sensor cover. Part of the hole is bored out for a snug fit with the motor, leaving about the last one-quarter inch of depth for the pump casing itself. The rotor is a single blade of plastic or metal, about one-eighth inch wide and seven-sixteenths long, cemented onto the motor shaft. The pump inlet, which cannot be seen in Figure 6, is a one-quarter inch hole at the center of the casing, communicating with the one-eighth inch hole, which can just be seen at the center of the rate sensor cover.

The larger motors will require somewhat different pump construction than that shown in Figure 6, and you may want to mount the pump separately from the rate sensor. Always draw air through the rate sensor from behind the thermistors, rather than forcing it into the outer end of the jet nozzle. The return path for the air is through the case in which the rate sensor is mounted, so this case must be reasonably tight to avoid any internal drafts that would disturb the flow of air being drawn into the jet nozzle. Do not connect flexible tubing to the nozzle inlet, as subsequent shifting of the position of the tubing will change the mechanical zero adjustment.

If possible, mount the rate sensor in the same case with the rest of the wing leveler. If this is not practical, mount the rate sensor in its own case and connect it to the wing leveler with shielded leads. Both cases should be of metal and grounded, to avoid interference from your radio transmitter; or, if plastic, should be lined with tinfoil, and the foil grounded.

Tilt Angle

The angle at which the rate sensor is mounted determines the damping factor of the wing-leveler-airplane combination (Ref. 8). The sensor is mounted so that the nozzle is toward the front of the aircraft (blowing toward the rear) with the outer end of the nozzle tilted upward.

A tilt angle of 30 degrees from the horizontal is a good average for all the aircraft we have had any experience with. This angle is not critical, but if you want to optimize it for your bird, find some still air, trim up the wing leveler and set the gain for normal operation. Then put the aircraft into a steep bank and release the control stick. If the aircraft is sluggish in returning to level, you have too great a tilt angle or too little control authority; if it overshoots appreciably, the tilt angle is too small. A good wing leveler with adequate control authority and optimum tilt angle will return the airplane to level from a 50 degree bank in seven or eight seconds.

Instrument Cases

By judicious layout, the circuit of Figure 3 can be squeezed into one of those neat digital clock cases available from Radio Shack (270-285), as shown in Figure 7. Some of the corners have to be filed off the rate sensor, and the nozzle must protrude from a slot in the top of the case and be boxed in to avoid cockpit drafts. This makes a neat under-the-dash installation for KRs and other cozy-cockpit birds, and there appears to be enough room for the addition of a magnetic-heading-hold circuit board. The whole installation, including servo, weighs less than twelve ounces.

If you have a spare instrument hole in your panel, the layout used by Omnics (Ref. 5) is much less of a hassle to implement. They use a standard 2" x 3" x 5" aluminum case (RS 270-238) with a control panel to fit a 3 inch diameter hole mounted on the 2" x 3" end of the case. With the three-inch dimension vertical, there is plenty of head room for the rate sensor at a thirty degree or more tilt angle, and much more area for the circuit board.

I would recommend getting everything possible on a single printed circuit board. Every extra plug and connecting wire is a potential source of trouble in an airborne installation.

Be sure to fuse the whole business when you install it in your airplane — there's no point in burning up your whole airplane if something goes wrong. Use an automotive in-line fuse holder in the power lead if there is no room for a fuse in the device itself. A one-and-one-half or two amp fuse should take care of the wing leveler and one servo.

Control Surface Matters

All the test flights of the electro-fluidic autopilot that we have made here at Langley have used vacuum-powered servos, acting directly on the aircraft's aileron linkages, so I have no direct flight experience with the R C servos. However, a few general comments on this area should not be out of order.

The type of installation you can use will depend, primarily, on the amount of friction in your aileron linkages. If your controls are well designed and well rigged, so that there is very little friction, you can use a servo-actuated tab on one aileron, or even couple the R C servo to the aileron linkage through a light spring, if control forces are low enough.

If you have a sticky control system, you will need a "hard" servo on the aileron controls, and this is a bit beyond the scope of this article. Another way out for the high-friction systems is a "separate-surface" arrangement as pioneered by Professor Jan Roskam (Ref. 9) and used by Don Hewes in one of the early homebuilt installations (Ref. 10, 11). The latter used a small, auxiliary control surface on each wing tip, positioned by individual R C servos, and developed a criteria for sizing these surfaces. If you take this route and use two servos, you will need a separate five-volt regulator for each, and a bit more fuse capacity.

If you opt for servoed tabs, you may have to experiment a bit on tab sizing. The text book (Ref. 12) says: "Trimming tabs have a chord varying from 5 to 10 percent of the movable surface chord. The aspect ratio should be as high as possible, varying usually from as low as 5 to as high as 20."

You will probably want to go a bit larger than this to get enough control authority to handle rough-air flying. You don't want the things to be so big that those puny little servos can't handle them, and you don't want to have to fight the control stick too hard all the way home if something goes wrong, and the tab goes full over. If you can get about 25 percent of full aileron deflection for full tab deflection, you should have plenty of aileron effectiveness for good wing-leveler performance.

There have been several successful installations in T-18s, and the tab on the one I saw appeared to be about two by eight inches.

Bob Paschal of Orlando, Florida, told me about a very satisfactory installation in his Ercoupe. He ended up with a three by twelve inch sheet metal tab hinged near the outboard end of one aileron. This gave him plenty of control authority for rough air flying, and was still small enough for the R C servo to handle.

Please be very careful with any of these control installations. If you tie into the main aileron control linkage, be very sure that whatever you add does not limit your maximum control deflection; will not apply control torques that are difficult to overpower; and that any conceivable breakdown will not jam up your aileron controls.

If you use a servoed trim tab, bear in mind that a failure in the linkage, allowing the tab to float free, **may well lead to a fatal flutter condition.** So do a good job on the linkage, and inspect it often. Also, remember that if you mount the servo in the aileron structure, you may upset its mass balance.

Norm Smith, who makes those quality electric trim systems (Ref. 13), puts a "mouse trap" spring in the hinge of his trim tab, so that if the linkage fails, the tab will snap up. You have to fight some out-of-trim on the way home, but he claims this is an effective flutter-preventer.

Well, that's about it for this time. As a parting thought, I should like to remind you that just because an airplane flies, does not necessarily mean that it's a good airplane — and the same thing applies to an autopilot. As Robert Townsend says in Up The Organization, "If you can't do it excellently, don't do it at all." Your life just might depend on it.

References

1. Garner, D.: The Saga of the Plastic Autopilot, *SPORT AVIATION*, March, 1974.
2. Garner, D.: Construction Notes on Electro-Fluidic Wing Levelers, *SPORT AVIATION*, June, 1978.
3. Moore, A. G.: Angular Movement Sensing Device, U. S. Patent 3,943,763, March 17, 1970.
4. Coughlin, Robert F. and Driscoll, Frederick F.: Operational Amplifiers and Linear Integrated Circuits, Prentice Hall 1977.
5. Omnics Corporation, P. O. Box 371, Ansonville, NC 28007.
6. Wilodru, Inc., Box 16, Raemont Road, Granite Springs, NY 10527.
7. Micro Mo Electronics, Inc., 3691 Lee Road, Cleveland, OH 44120, (216) 921-1131.
8. Garner, H. Douglas and Poole, Harold E.: Development and Flight Tests of a Gyro-Less Wing Leveler and Directional Autopilot. NASA TN D-7460, 1974.
9. Roskam, J., Barber, M. R., and Loschke, P.C.: Separate Surfaces For Automatic Flight Controls, SAE paper 730304 April, 1973.
10. Hewes, Don: Preliminary Report — Development of a Poor Man's VFR Autopilot, *SPORT AVIATION* , May, 1978.
11. Hewes, Donald E.: Update on the Poor Man's Autopilot, *SPORT AVIATION*, March, 1979.
12. Teichmann, Fredrick K.: Airplane Design Manual, Pitman, 1958.
13. Aero-Trim, Inc., 1130 102 St., Bay Harbor, FL 33154.

This four place KZ 7 Lark is very interesting in that with 125 hp (Continental), it is similar in size, weight and overall configuration with the early Piper PA-20 Pacers. The wing of the KZ 7 is much more sophisticated, however, what with its full span fixed slots, big mass balanced ailerons and flaps. Top and cruise speeds are similar . . . but the KZ's 34 mph landing speed puts it in a class of its own.